Open Science Grid

Alan Sill

Texas Tech University/
Fermilab

*Consortial International Workshop
on Computational Physics
Hsinchu, Taiwan, Dec. 2 - 4, 2004*

Elements of Global Computing

(Personal point of view)

- It should be easy to adopt
  - The tools, concepts and workflow should look familiar to the scientists and users.
- It should work for everyone
  - E.g in particle physics, entire collaboration built the detector => entire collaboration analyzes the data
    - *Keep user environment simple!*
- It should add or create additional resources
  - There should be some distinct advantage obtained from the computing being globally distributed.
- How can we go from these general principles to computing that is *truly* global and pervasive?

The Open Science Grid

- Not a start from scratch!
- Not a *project*, but a *consortium*!
- Begun as an extension from the earlier Grid2003 effort.
  - Became "Grid3" as it achieved success beyond '03
- Explicitly multi-disciplinary.
- Explicitly trying to build a blueprint and procedure:
  - "What are the general best practices for grids extracted from previous efforts, and how can we best create templates and methods to implement them?"

Goals of the OSG:

- Provide a framework into which grids can be integrated to allow user communities to benefit from the best available expertise in grid technology
- Provide an operational model while not restricting users to be from a single user community
- Share best practices in grid security, operations, use policy, technical architecture, monitoring, and data handling/storage for shared science operations
- Be prepared to generalize to very large numbers and complexity of resource deployment.

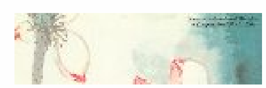

OSG Structure (12/2004)

Organized around *Technical Groups* and *Activities* to connect *Virtual Organizations* with *Service Elements* and *Resource Providers*:

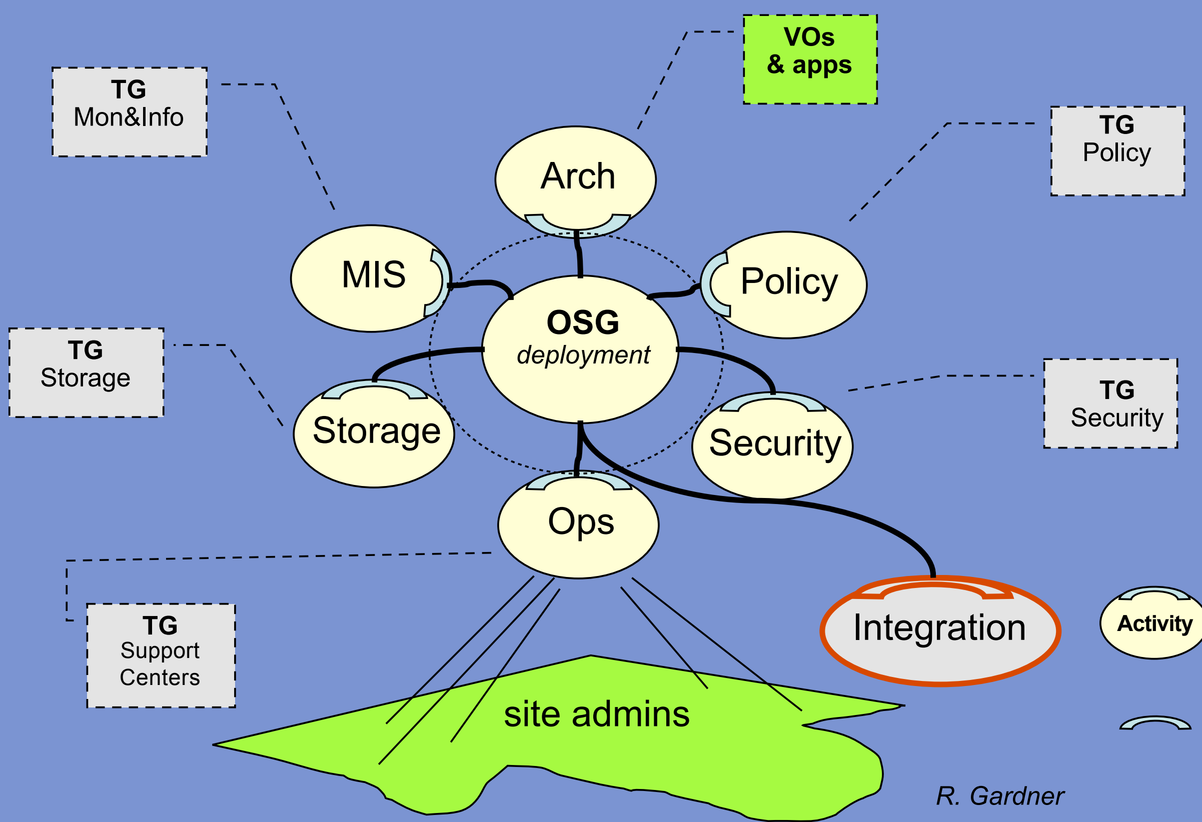

OSG is VO-based:

- Applications (VO-based) interact with Facilities through infrastructure:

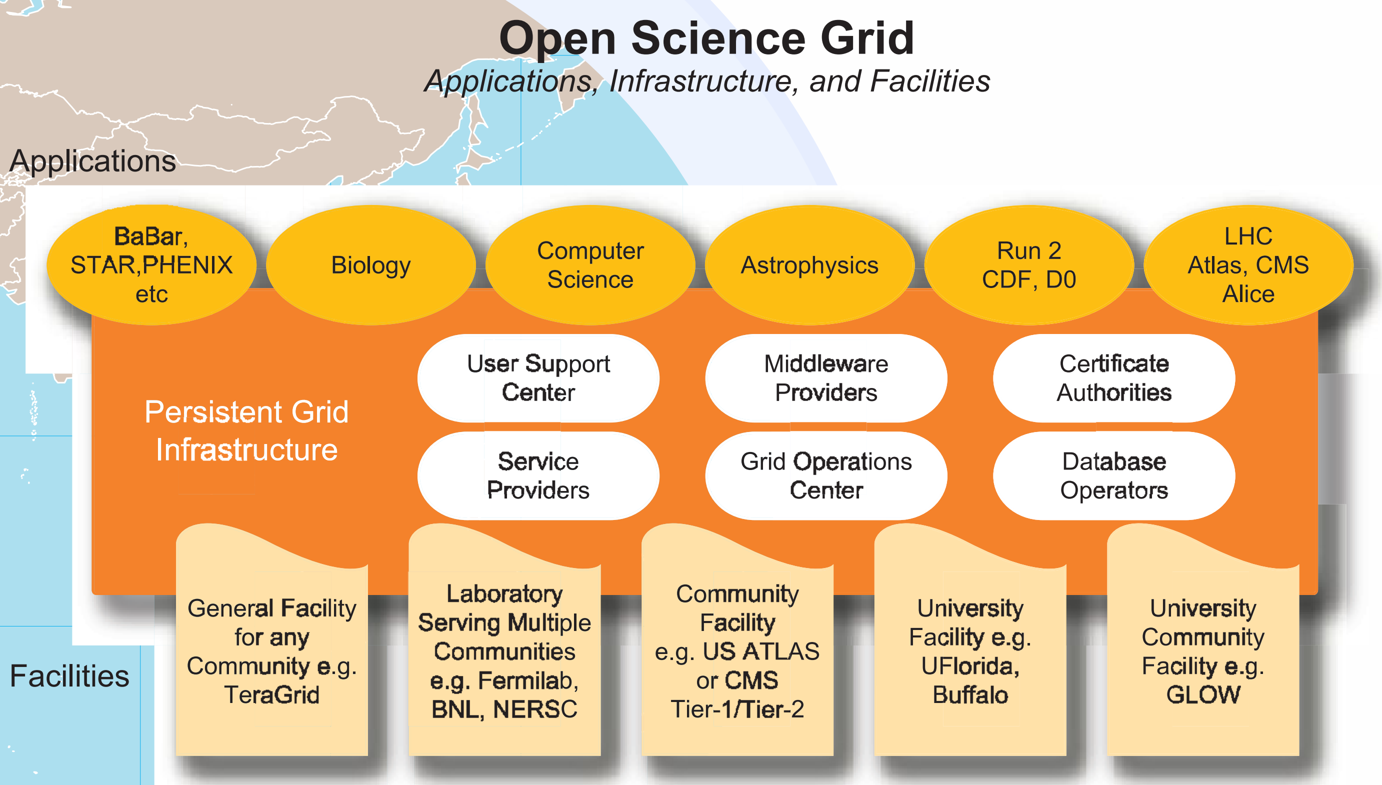

OSG Governance

- Open to contributions based on participation:

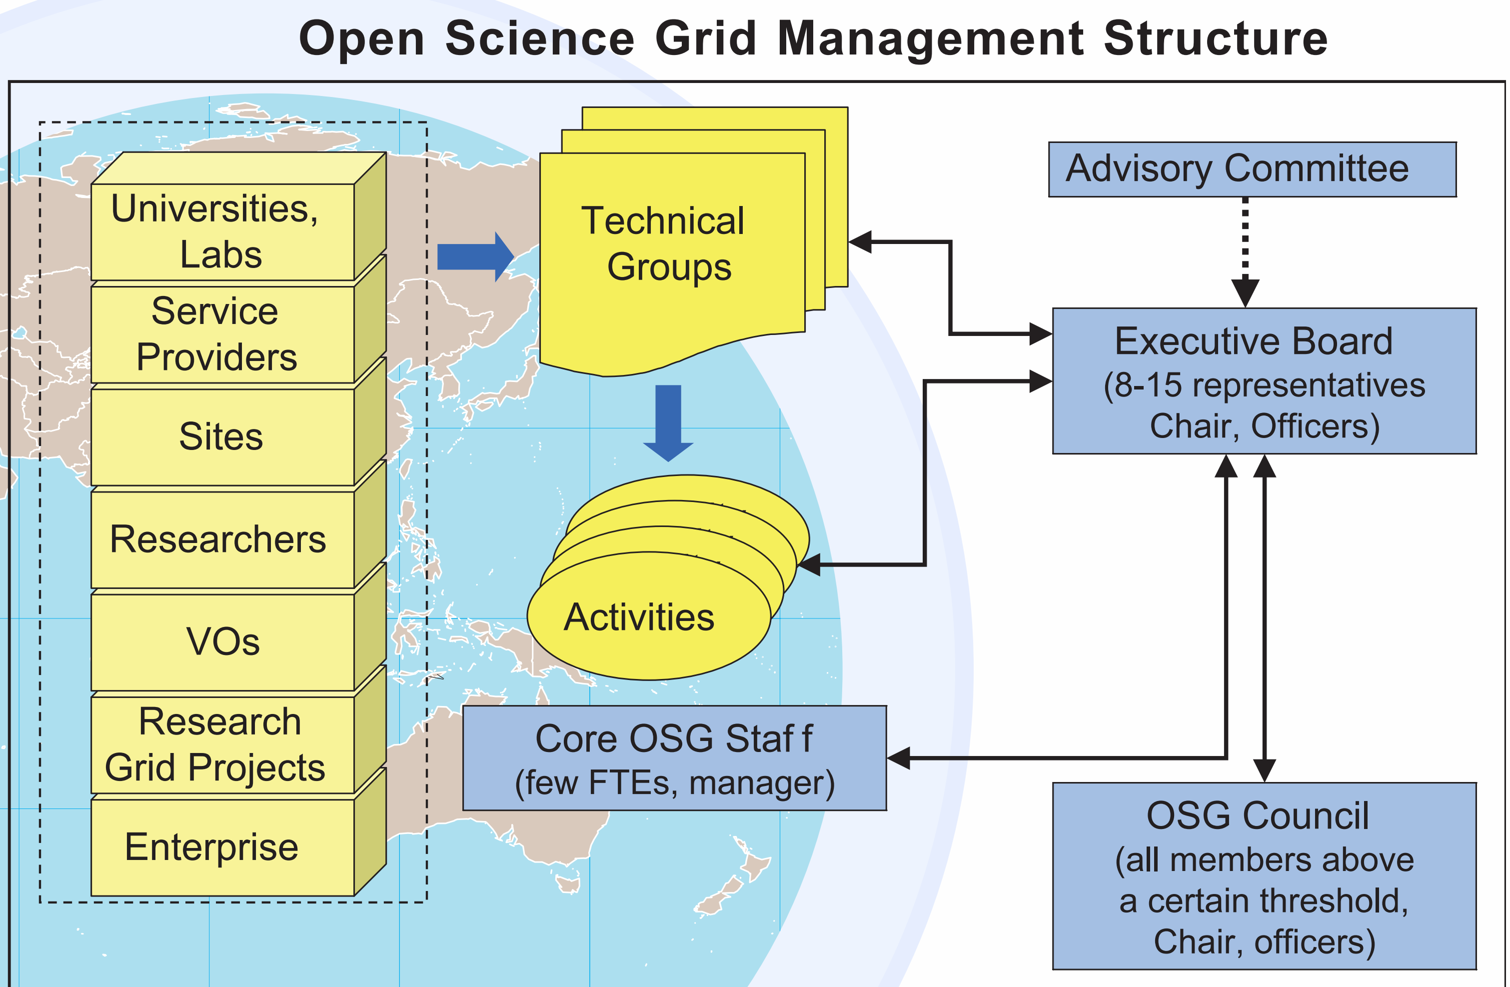

Related Goals:

- Provide a US national infrastructure for science grids in a way that is *explicitly inter-operational* with other national and global efforts.
- Produce templates and recommendations in terms of "best-practice" documents, and a blueprint to guide Virtual Organizations and Resource Providers in mutually beneficial arrangements for usage.
- Share technology among user communities.
- Make useful contact with and extract experience and information from other world-wide efforts.

Grid3 Deployment Snapshot

Jan. 2004

- Community of ~100 users
- Grown to 100% utilization

- 30 sites, multi-VO
- shared resources
- ~3000 CPUs

Sep. 2004

Example: Grid3/iVDGL Operations matrix

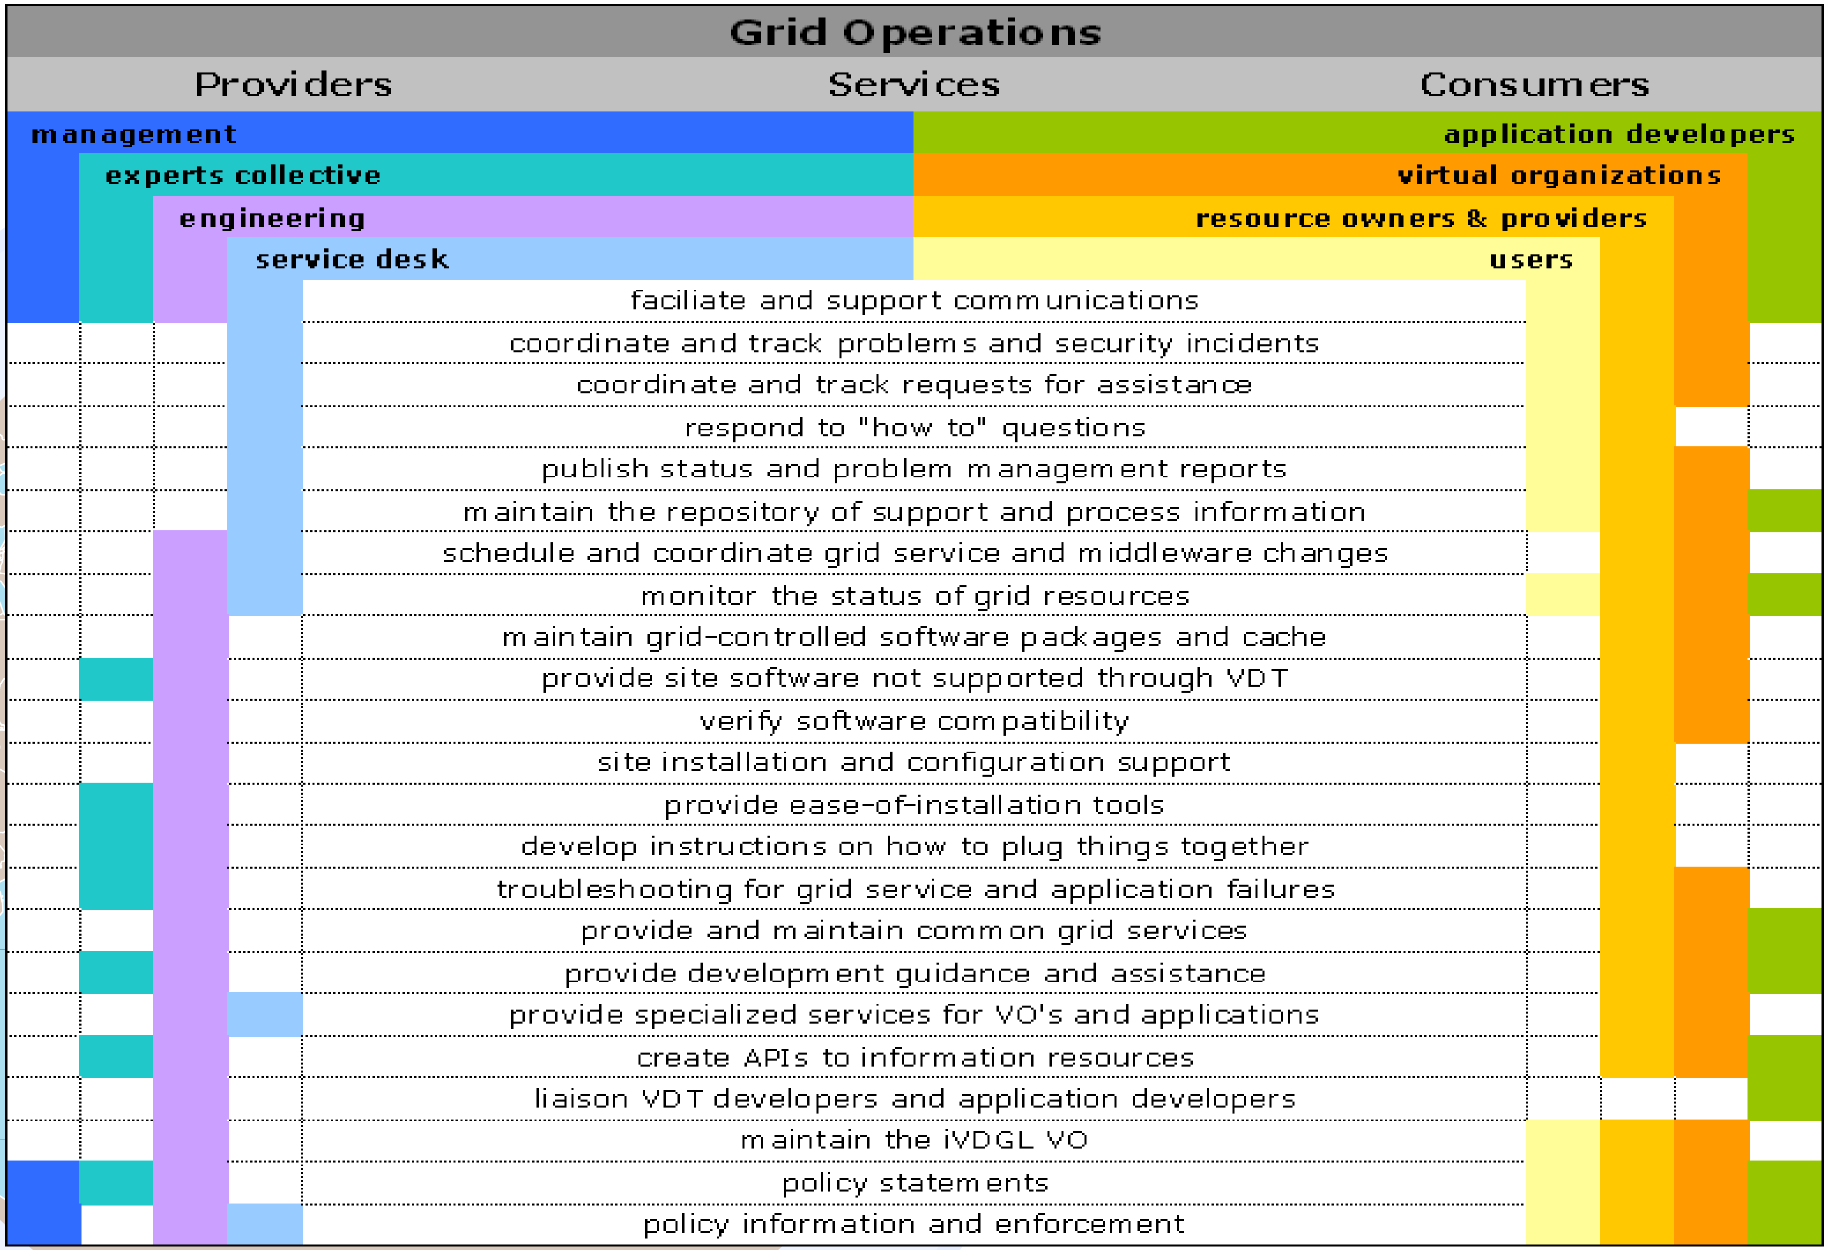

OSG is not just for US projects!

- Example: biggest users of Grid3 up to now have been LHC experiments:

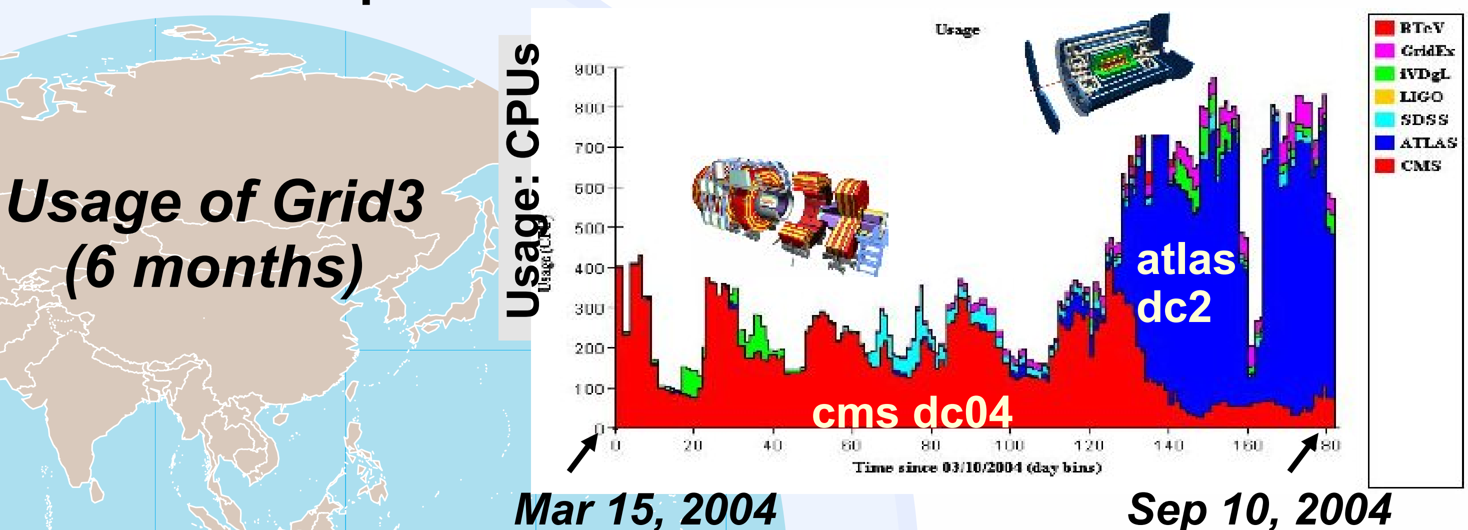

- Similarly CDF, Dzero based in the US are already highly international.
- Other fields (computational chemistry, biology, etc.) also have a strong non-US basis.

Security and Incident Response

- One of our strongest efforts.
- As with others, based on an Activity advised by a Technical Group.
- Strong overlap and interaction with EGEE.
- Documents downloadable and written so far:

  http://www.opensciencegrid.org/techgroups/security/#atwork

- Other related technical groups: Policy, Governance
- Example of useful outcomes: "Privilege Project" (http://computing.fnal.gov/docs/products/voprivilege/)

Current Grid3 Multi-VO Security Model

- DOEGrids Certificate Authority
- PPDG and iVDGL RA with VO or site sponsorship
- Automated multi-VO authorization, using VOMS
- Each VO manages a service and its members
- Each Grid3 site is able to generate a Globus gridmap file with an authenticated SOAP query to each VO service
- Site-specific adjustments or mappings

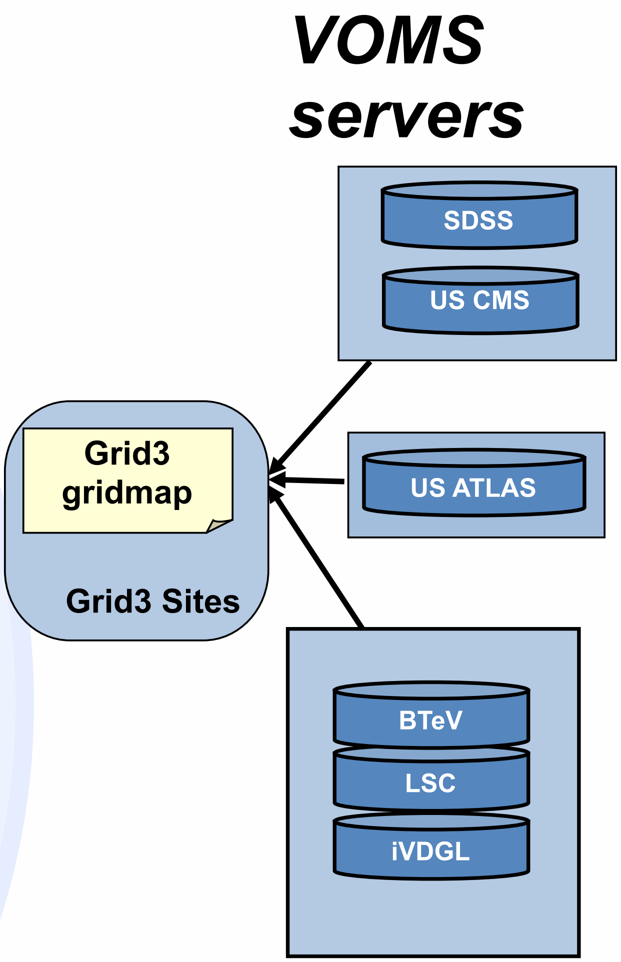

Example extension: Privilege Project

- Extend above model to include *privileges* and *roles* to allow people to come in with different VO-based flags for different purposes
- Allows combining of VO resources for multiple purposes
- Example of technical capability that allows policy decisions.

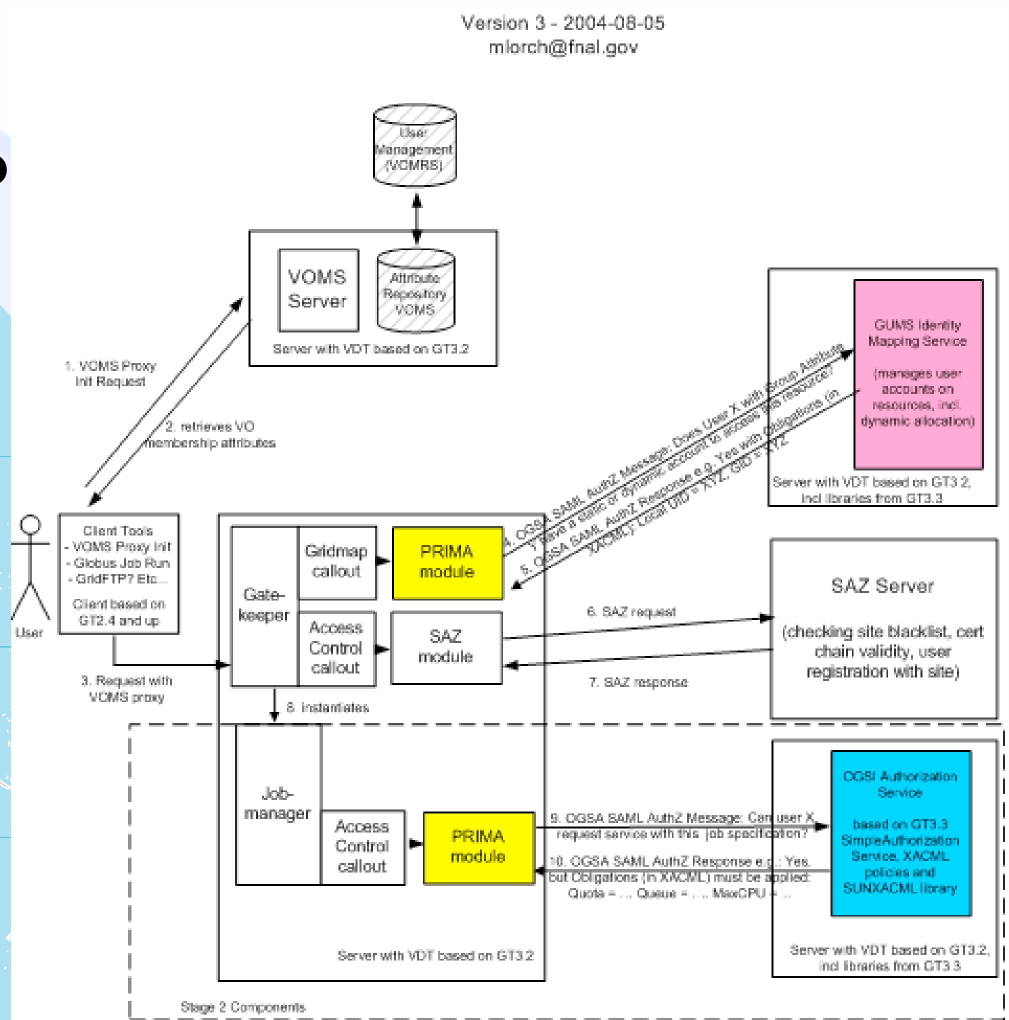

Examples of Principles in the Blueprint

(Dec 1st 2004, I Open Science Grid - Operations Workshop):

- OSG will provide baseline services and a reference implementation.
- Use of other services will be allowed.
- All services should support the ability to function and operate in the local environment when disconnected from the OSG environment.
- Users are not required to interact directly with resource providers.
- The requirements for participating in the OSG infrastructure should promote inclusive participation both horizontally (across a wide variety of scientific disciplines) and vertically (from small organizations like high schools to large ones like National Laboratories).
- The OSG architecture will follow the principles of symmetry and recursion.
- The OSG infrastructure must always include a phased deployment, with the phase in production having a clear operations model adequate to the provision of production-quality service.

Run II CDF Central Computing

- Example of resources not yet fully integrated onto the grid.
- CDF, Dzero: developed in 1999–2002 to respond to each experiment's greatly increased need for computational and data handling resources to deal with Run II.
  - Greatly increased cpu power & data to physicists (e.g. CDF: presently 300 TB data + 1500 cpus).
- Among the first large-scale cluster approaches to general user analysis computing for ubiquitous use (as opposed to special farms for production of canned jobs).
- Replaced Symmetrical Multi-Processing (SMP) approach with inexpensive collections of commodity Linux-based computing and file server systems.

Example: CDF Global Analysis (2004)

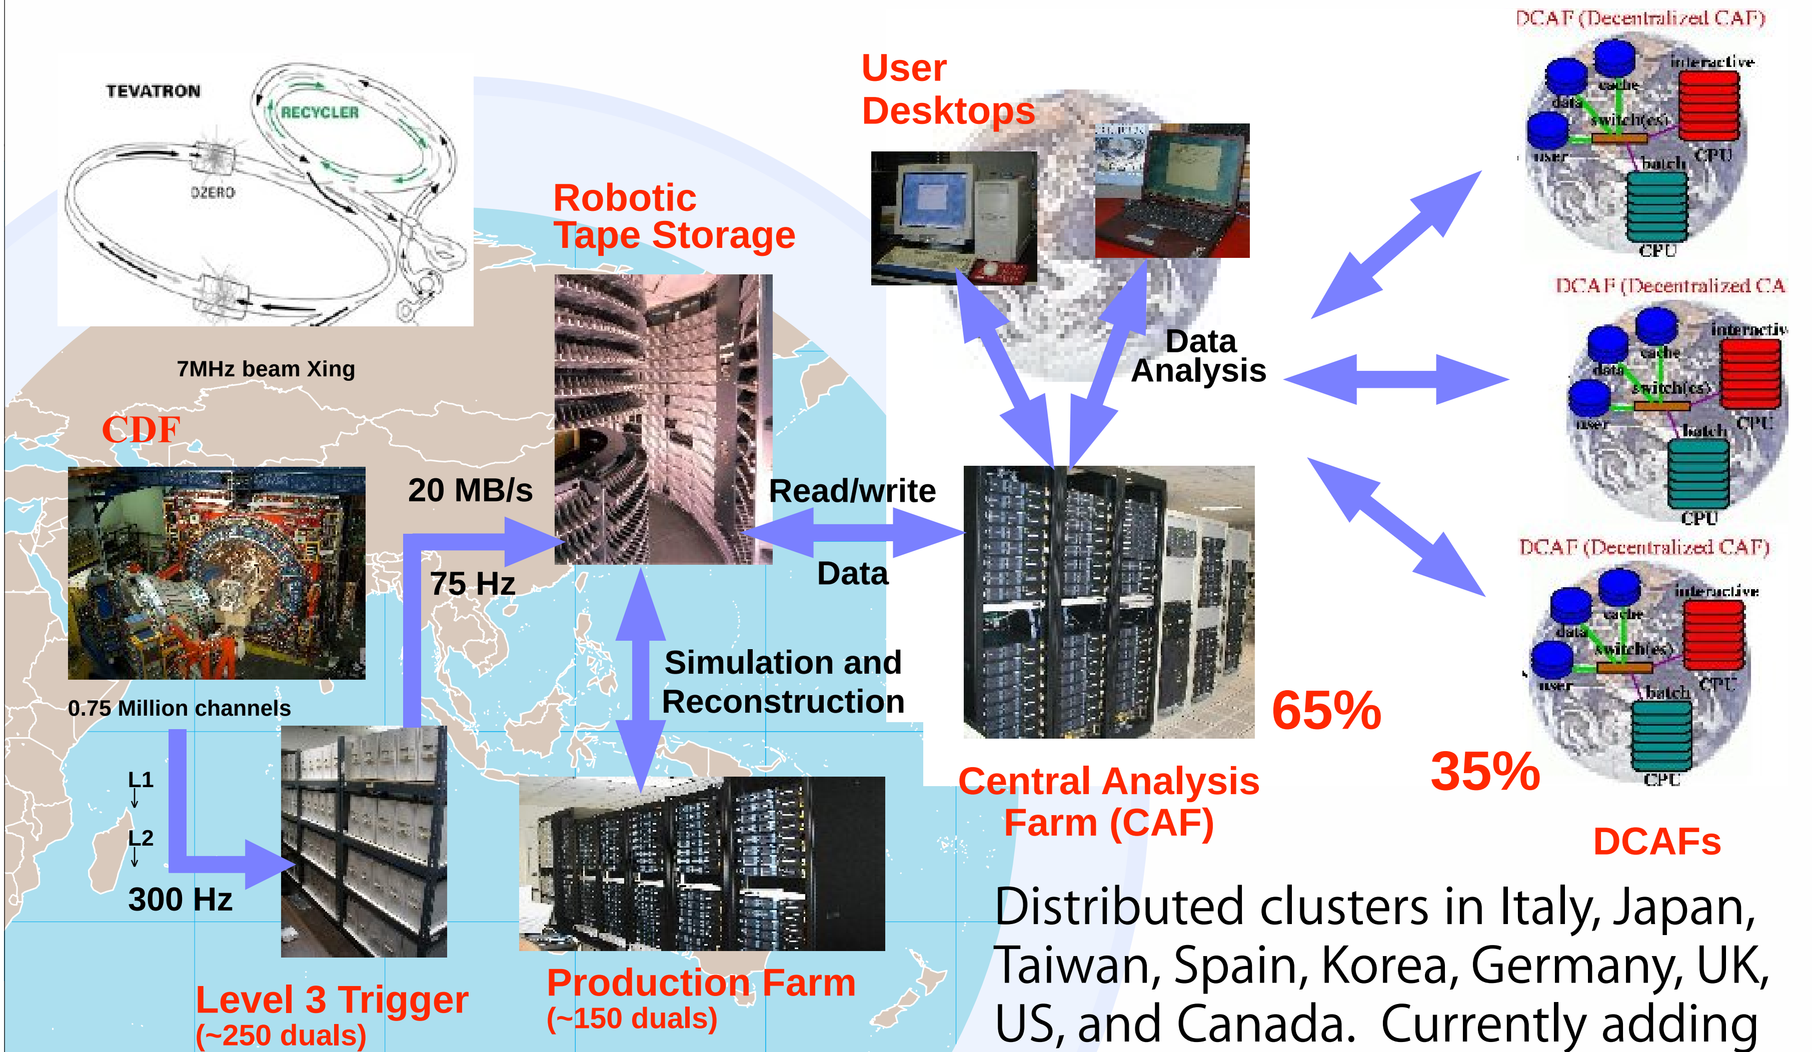

Distributed clusters in Italy, Japan, Taiwan, Spain, Korea, Germany, UK, US, and Canada. Currently adding Grid3, LCG resources to this picture.

Current CDF Dedicated Resources:

Current Resources [*]			As of Sep. 24, 2004
Cluster Name and Home Page	Monitoring and Direct Information Links	CPU (GHz)	Disk space (TBytes)
Original FNAL CAF	queues, user history, ganglia, sam station, consumption	1200	300
FNAL CondorCAF (Fermilab)	queues, user history, analyze, ganglia, sam station, consumption	2000	(shared w/CAF)
CNAFCAF (Bologna, Italy)	queues, user history, resources, network, sam station, datasets, consumption	300	7.5
KORCAF (KNU, Korea)	queues, user history, ganglia, sam station, datasets, consumption	120	0.6
ASCAF (Academia Sinica, Taiwan)	queues, user history, ganglia, sam station, datasets, consumption	134	3.0
SDSC CondorCAF (San Diego)	queues, user history, analyze, ganglia, sam station, datasets, consumption	280	4.0
HEXCAF (Rutgers)	queues, cpu, sam station, datasets, consumption	100	4.0
TORCAF2 (Toronto CDF)	queues, ganglia, disk status, sam station, datasets, consumption	576	10
JPCAF (Tsukuba, Japan)	queues, user history, ganglia, sam station, datasets, consumption	152	5.0
CANCAF (Cantabria, Spain)	queues, user history, ganglia, sam station	52	1.5
MIT (Boston, USA) (MC only)	queues	110	2.0
* *(Counts only resources openly available to all CDF users)*	*Current Totals [*]:*	5024	337.5

Moving from dedicated resources to the OSG:

- CDF, Dzero good examples of organizations that have achieved early success through deployment of dedicated resources.
- To go further, need to move to using these resources in a more global way, add grid capability and gain the ability to interact with other grid-enabled tools.
- Overhead of maintaining a separate infrastructure for each such experiment or use is too great.
  - Software development.
  - Strain on experts, retraining of new experts, etc.
  - Transience and pace of development of knowledge.
  - => Move to new paradigm based on globally developed tools.

Federation of Grids/Interoperability

- "Grid of grids" an inevitable consequence?
- Use OSG to find out!

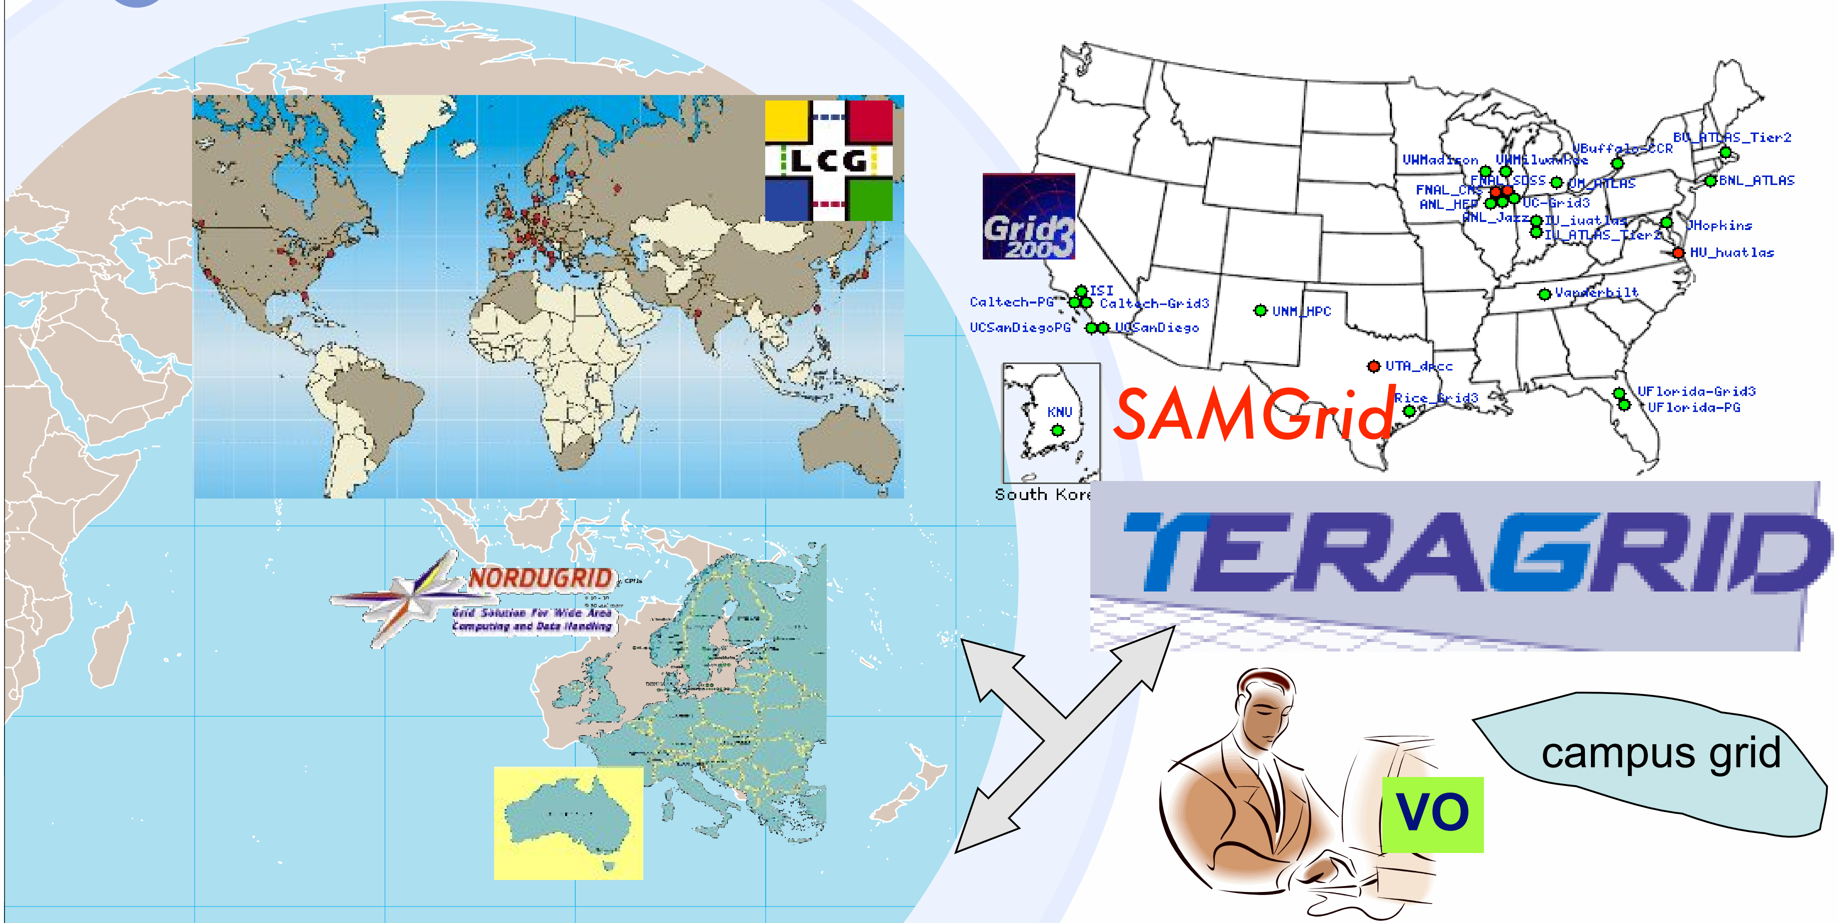

The Texas High Energy Grid (*THEGrid*)

- Example of a Regional Grid affiliated with OSG.
- Applications: HEP, comp.chem, biogrid, cancer therapy, etc.

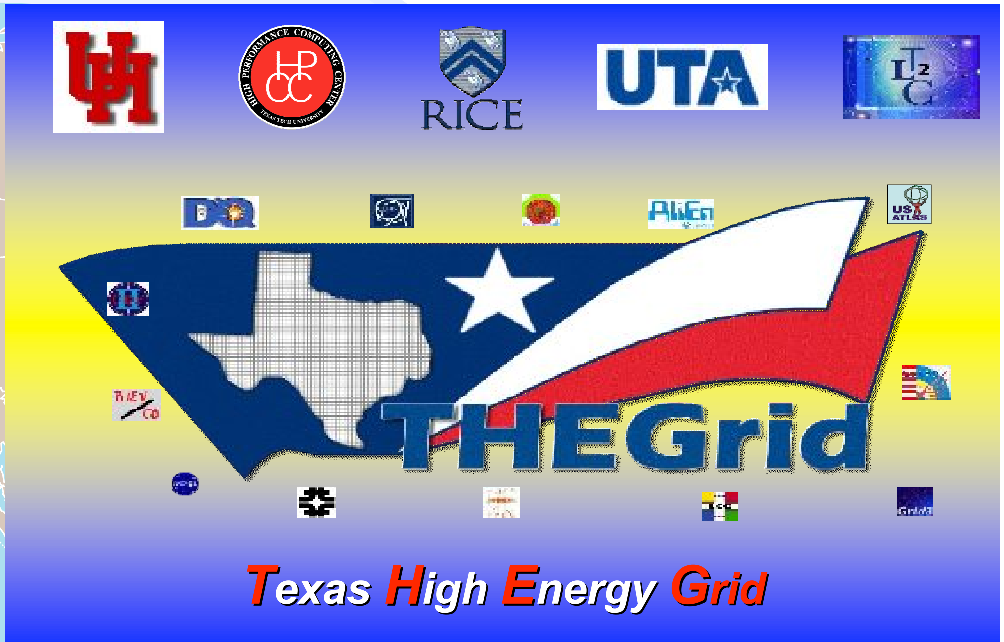

- Will participate in OSG development and apply Blueprint.
- Good model for development and for use in Asia - can serve multiple VOs/users w/same grid!

Conclusions

- Future of high performance computing development is highly generalized access to multiple resources organized along discipline-specific usage pattern.
- OSG provides a template for the general design of such resources, not just a single grid.
- Open to participation and new ideas.
- Will provide a specific implementation as a common resource, but organized to be cooperative and cross-disciplinary for all of science.
- Please join and get involved!

Where to go for more information:

DOE Office of Science

NSF

Open Science Grid

http://opensciencegrid.org